

# Graph Theoretic Investigations on Inefficiencies in Network Models

Pietro Cenciarelli, Daniele Gorla and Ivano Salvo  
Sapienza University of Rome, Dpt. of Computer Science  
{cenciarelli,gorla,salvo}@di.uniroma1.it


**Abstract.** We consider network models where information items flow from a source to a sink node. We start with a model where routing is constrained by energy available on nodes in finite supply (like in Smartdust) and efficiency is related to energy consumption. We characterize graph topologies ensuring that every saturating flow under every energy-to-node assignment is maximum and provide a polynomial-time algorithm for checking this property. We then consider the standard flow networks with capacity on edges, where again efficiency is related to maximality of saturating flows, and a traffic model for selfish routing, where efficiency is related to latency at a Wardrop equilibrium. Finally, we show that all these forms of inefficiency yield different classes of graphs (apart from the acyclic case, where the last two forms generate the same class). Interestingly, in all cases inefficient graphs can be made efficient by removing edges; this resembles a well-known phenomenon, called *Braess's paradox*.

## 1 Introduction

Through the years, several formal models have emerged for studying network design in terms of network traffic, protocols, energy consumption, and so on (see [4,15,18,12], just to cite a few). Our investigation of energy efficiency and load balancing of multi-hop communication in ad-hoc networks started in [6], where we considered a simple model, called *depletable channels*, in which networks are oriented graphs with nodes equipped with a natural number representing depletable charge, as in Smartdust [22].

To better understand our model, consider a scenario in which four communication devices are located in a landscape (Fig. 1(a)). The devices may have different communication radii, e.g. because of geographical reasons or because of their settings (a larger radius requires more energy to perform a communication). Thus, node reachability is not symmetric: in our example,  $s$  can reach (i.e., send information to)  $u$  and  $v$ , but not vice versa. Such a scenario can be very naturally modeled via a directed graph, where vertices correspond to communication devices and an edge is placed between  $x$  and  $y$  if  $y$  falls within the communication circle centered in  $x$ . In our example, the resulting graph is given in Fig. 1(b). To simplify reasoning, we will always assume that information flows from a single source (a device that does not receive information from anyone else –  $s$  in our example) to a single destination (a device that does not send information to anyone else –  $t$  in our example).

Devices have a depletable amount of energy that is consumed throughout their life; moreover, they are just information forwarders, so they can only send or receive information. Usually, a device consumes less energy when it receives than when it sends; moreover, energy consumption is proportional to the length of the message exchanged.



**Fig. 1.** (a) Four communication devices with their communication radius: (b) the associated directed graph.

This is modeled in our setting by slotting both information and energy, and by assuming that sending one information unit consumes one energy unit. By contrast, receiving information does not lead to any consumption. Furthermore, energy can only decrease during the life of the system; so, we do not model any form of recharge. This implies that every system has a lifetime: after sending a certain amount of information from  $s$  to  $t$ , all the intermediate nodes (or at least all those in a cut) will eventually die, due to lack of energy. The resulting system is then called *dead* and the flow of information leading to it will be called *inhibiting*.

Aimed at capturing an abstract notion of communication service provided by a network, in [6] we introduced an equivalence on networks that equates two networks if and only if they have identical maximum and minimum inhibiting flow (there, we also proved that this corresponds to what in the theory of concurrency is called *trace equivalence*). Networks in which the minimum inhibiting flow is strictly less than the maximum flow can be regarded as inefficient, because an arbitrary routing may cause unnecessary depletion of energy. As in [12], we consider flows not controlled by a planning entity and thus any legal flow can take place. We call *weak* those graph that, for some charge assignment to nodes, the minimum inhibiting flow is strictly less than the maximum flow. For example, the channel whose topology is given by the graph in Fig.1(b) and where  $u$  and  $v$  have charge 1 (usually, we assume that  $s$  and  $t$  have 'big' charge) can deliver at most 2 information items. However, it is also possible that a single unit of information inhibits the channel, if it takes the 'wrong' way (viz.,  $s u v t$ ). Thus, the graph in Fig.1(b) is weak.

As shown in this paper, network performance can be improved by removing edges belonging to  $st$ -paths passing at least twice through a minimal vertex separator (that is a minimal set of vertices whose removal disconnects the source from the sink). In our example, removing the edge  $u \rightarrow v$  leads to a non-weak graph. This somehow resembles a counterintuitive but well-known situation in the setting of selfish routing in traffic networks, called Braess's paradox [2,5]. Braess's paradox occurs when the equilibrium cost may be reduced by raising the cost of an edge or, equivalently, by removing such an edge. This form of inefficiency is known as *vulnerability* [17,14]. Indeed, the work presented here started with the empirical observation that all natural examples of weak (acyclic) graphs turned out to be vulnerable. As this paper shows, this is not incidental.

Both depletable channels and traffic networks have strong similarities with the standard model of flow networks [1]. In particular, depletable channels can be easily translated into standard (edge-capacitated) flow networks. Usually, this translation is painless, as it preserves all the standard notions in flow networks, such as maximum and minimum inhibiting flow. Interestingly, if we define as *edge-weak* those graphs that, for some assignment of capacities to edges, admit a non-maximum inhibiting flow, it turns out that weakness and edge-weakness *do not* coincide.

In this paper, we compare these three kinds of inefficiency from a graph-theoretical perspective. Our main contributions are:

1. In Section 2, we characterize *weak graphs* as those graphs with an *st*-path that passes twice through the same minimal vertex separator.
2. In Section 3, we then show that weakness can be checked in polynomial time.
3. In Section 4, we characterize *edge-weak graphs* as those graphs with an *st*-path that passes twice through the same cut-set (i.e., the set of edges that connect the two blocks of a cut), for a proper kind of cut that we call *connected*.
4. In Section 5 we characterize *vulnerable graphs* as those graphs containing a subgraph homeomorphic to Fig.1(b). This extends previous results of [14,7] to arbitrary directed graphs and completely solves a question left open in [17]. Indeed, in [14] the characterization holds only for undirected graphs, whereas in [7] the result holds only for a specific class of directed graphs (that they call *irredundant*).
5. Finally, in Section 6, we relate all the above mentioned class of graphs. In the general case, vulnerability implies edge-weakness. Moreover, vulnerable graphs always contain an acyclic weak subgraph. This suggests that the core reason that makes a graph vulnerable is indeed its weak subgraphs. If we restrict our attention to DAGs, weakness implies both vulnerability and edge-weakness. These two classes coincide and they coincide with the class of graphs that are not series-parallel [16].


Due to space constraints, all proofs are relegated to the Appendix.

## 2 Inefficiency in Depletable Channels: Weakness

A *depletable channel* is a graph  $G$  equipped with a function  $\eta$  associating with each node a non-negative integer representing its *depletable charge*. Throughout the paper, we only consider *directed simple st-graphs*, that are directed graphs without self-loops and parallel edges, with a fixed source vertex  $s$  and sink vertex  $t$ . When no confusion arises, we write just  $\eta$  for a channel  $(G, \eta)$  and call  $G$  its *underlying graph*.

Charges may change as result of information passing through the net. Each item passing through a node consumes one unit of the node's charge, thus leaving the channel in a state of lower energy. We also assume that the charge of the source and of the sink are always large enough to be irrelevant (say,  $\infty$  for simplicity).

An *st-path* (just path from now on) in a graph is a (possibly cyclic) directed walk from  $s$  to  $t$ . The set  $P(\eta)$  is the set of all paths of  $\eta$ . Since we only consider source-to-sink paths, we can assume that every vertex belongs at least to one of such paths. A *flow* for  $\eta$  is a function  $\phi : P(\eta) \rightarrow \mathbb{N}$  such that  $\phi(v) \leq \eta(v)$ , for every  $v \in V$ . Here,  $\phi(v)$  denotes the amount of  $v$ 's charge consumed by  $\phi$ , i.e.,  $\phi(v) = \sum_{p \in P(\eta)} r_{vp} \cdot \phi(p)$ , where  $r_{vp}$  is the number of times in which node  $v$  is repeated in the path  $p$  (0, if  $v \notin p$ ). With respect to standard flow models, we admit a positive flow over cycles because in our



**Fig. 2.** Two channels

model each node knows its neighborhood only. Without global information about the net, routing can easily lead to cyclic paths.

The *value* of  $\phi$  is  $\sum_{p \in P(\eta)} \phi(p)$ . We denote by  $\max_{\eta}$  the maximum value of a flow for  $\eta$ . We call  $\eta$  a *dead* channel if  $\max_{\eta} = 0$ . The residual of a channel  $(G, \eta)$  after a flow  $\phi$  is a channel with the same underlying graph  $G$  and with a charge function  $\eta'$  such that  $\eta'(v) = \eta(v) - \phi(v)$ , for every vertex  $v$ . A flow  $\phi$  is said to *inhibit*  $\eta$  if the residual of  $\eta$  after  $\phi$  is dead. We denote by  $\min_{\eta}$  the smallest value of an inhibiting flow in  $\eta$ .

The form of inefficiency in this network model is given by the following definition.

**Definition 1.** A graph  $G$  is weak if  $\min_{\eta} \neq \max_{\eta}$ , for some channel  $\eta$  whose underlying graph is  $G$ .

Two typical examples of channels are depicted in Fig. 2, where we depict the charge of a node in place of its name and use 'o' to denote a large-enough charge. Here, channel (1) is not weak, since every routing of messages always ensures the delivery of 2 information units (remember that every information unit consumes one charge unit of every traversed node). By contrast, channel (2) is weak because there is also a inhibiting flow that only delivers 1 unit (by sending the information unit along the path  $o \rightarrow 2 \rightarrow 1 \rightarrow 2 \rightarrow o$ ).

Weakness can be characterized in terms of the existence of a walk that passes at least twice through a *minimal st-separator* (or *mvs*, for minimal vertex separator, assuming  $s$  and  $t$  fixed). An mvs [11] is a minimal set of vertices whose removal disconnects  $s$  and  $t$ .


**Theorem 1.** A graph is weak if and only if there exists an mvs  $T$  and a directed walk  $a \rightsquigarrow b$  with  $a, b \in T$ .

One direction is proved by assigning charge 1 to all nodes in  $T$  and a big value to the remaining ones; we then exploit  $a \rightsquigarrow b$  to saturate  $T$  with a flow of value smaller than  $\max_{\eta}$ . For the converse, the existence of an inhibiting flow of value smaller than the maximum implies existence of an mvs saturated by this flow; this can only happen if there is a path passing twice through the mvs.


### 3 Checking Weakness

In [6], we proved that, given a graph and a charge  $\eta$  to its nodes, it is NP-hard to determine whether  $\min_{\eta} \neq \max_{\eta}$ . By contrast, we show in the following that checking weakness is a polynomial-time problem.

Stemming from Theorem 1, a trivial algorithm to determine if a graph is weak is to generate all mvs's and, for each of them, check if there exists a walk that touches the mvs twice. Unfortunately, the number of mvs's in a graph can be exponential in  $|V|$  [13]. However, as we now show, it is enough to examine at most  $|V|^2$  mvs's.



**Fig. 3.** A graph with a complete chain not containing a specific critical node



**Fig. 4.** A graph with a complete chain not containing any critical mvs

From now on, we will use the following terminology, inspired by Theorem 1. A walk  $a \rightsquigarrow b$  is *critical* if there exists an mvs  $T$  such that  $a, b \in T$ . In such a case, we call  $b$  a *critical node* and  $T$  a *critical* (or even *b-critical*) mvs.

We say that a node  $u$  is *covered* by a set of nodes  $A$ , notation  $u \sqsubseteq A$ , if all walks starting from  $u$  to the sink  $t$  contain at least a node  $v \in A$ . A set of nodes  $A$  is *covered* by a set of nodes  $A'$ , notation  $A \sqsubseteq A'$ , if  $u \sqsubseteq A'$ , for all  $u \in A$ . A set of nodes  $A$  *precedes* a node  $u$ , notation  $A \leq u$ , if all walks starting from the source  $s$  to the node  $u$  contain at least a node  $v \in A$ . A set of nodes  $A$  *precedes* a set of nodes  $A'$ , notation  $A \leq A'$ , if  $A \leq u$ , for all  $u \in A'$ . Since mvs's are minimal sets of *st*-separators, the following claim can be easily proved.

**Lemma 1.** *If  $A \leq a$  or  $a \sqsubseteq A$ , then  $A \cup \{a\} \not\sqsubseteq T$ , for every mvs  $T$ .*

Moreover, it is well-known [9] that the set of mvs's of a graph is partially ordered w.r.t.  $\sqsubseteq$ , with minimum element  $\{s\}$  and maximum element  $\{t\}$ . A sequence of mvs's  $T_0, T_1, \dots, T_n$  is a *chain* if, for all  $i$  ( $0 \leq i < n$ ), we have  $T_i \sqsubset T_{i+1}$ ; it is *complete* if  $T_0 = \{s\}$ ,  $T_n = \{t\}$  and  $T_i \sqsubseteq T \sqsubseteq T_{i+1}$  implies  $T = T_i$  or  $T = T_{i+1}$ , for every mvs  $T$ .

From [19], it follows that a complete chain contains at most  $|V|$  mvs's. Following the idea of finding a critical mvs and a critical walk by examining a complete chain of mvs's, we first observe that a specific critical node may never appear in a complete chain. To see this, consider the graph depicted in Fig. 3. In the mvs  $\{a, b, c\}$ ,  $b$  is a critical node, but it does not appear in the complete chain  $\{s\}, \{a, p\}, \{u, p\}, \{u, c\}, \{t\}$  that, however, contains the critical node  $u$ . This is not incidental, as the following theorem states.

**Theorem 2.** *If the graph is weak, then in every complete chain of mvs's  $T_0, T_1, \dots, T_n$  there exists at least a  $T_i$  that contains a critical node.*

Unfortunately, Theorem 2 is not enough to conclude that any complete chain of mvs's in a weak graph contains at least a critical mvs. To see this, let us consider the graph of Fig. 4: it is weak, as the critical mvs  $\{a, b\}$  testifies. However, if we consider the complete chain  $C' = \{s\}, \{a, p\}, \{a_1, a_2, p\}, \{a_2, b\}, \{t\}$ , we cannot find any critical mvs;  $C'$  just contains the non-critical mvs  $\{a_2, b\}$  that contains the critical node  $b$ .

Theorem 3 ensures that we can check if any node  $b$  is critical (and hence check the existence of a *b-critical* mvs) by considering any *b-minimal* mvs, i.e., an mvs  $T$  such that  $b \in T$  and  $b \notin T'$ , for every mvs  $T' \sqsubset T$ .

**Theorem 3.** *If  $b$  is a critical node, then all  $b$ -minimal mvs's are  $b$ -critical.*

In the example of Fig. 4, it suffices to consider the mvs  $\{a_2, b\}$ : its *b*-minimal predecessor (w.r.t.  $\sqsubset$ ) is  $\{a, b\}$ , that is *b-critical*.

---

**Algorithm 1** Checking Weakness

---

**Input:** A directed  $st$ -graph  $G = (V, E, s, t)$ ,  $s$  is the source and  $t$  is the sink

**function**  $weakOrNotWeak(G)$

```
1.  $T \leftarrow \{s\};$ 
2. while  $T \neq \{t\}$  do
3.    $T' \leftarrow immediateMvsRight(T)$ 
4.   if  $T'$  is critical then
5.     return WEAK
6.   for all  $b \in T' \setminus T$  do
7.      $T^* = minimalMvs(T', b)$ 
8.     if  $T^*$  is critical then
9.       return WEAK
10.   $T \leftarrow T'$ 
11. return NOT WEAK
```

---

---

**Algorithm 2** A  $b$ -minimal mvs smaller (w.r.t.  $\sqsubset$ ) than  $T$ 

---

**Input:** An mvs  $T$ , a node  $b \in T$

**function**  $minimalMvs(T, b)$

```
1.  $A \leftarrow \{u \in T \mid (T \cup pred(u)) \setminus \{b\} \not\sqsubseteq b\}$ 
2. while  $A \neq \emptyset$  do
3.   choose  $u \in A$ 
4.    $T \leftarrow T_u$ 
5.    $A \leftarrow \{u \in T \mid (T \cup pred(u)) \setminus \{b\} \not\sqsubseteq b\}$ 
6. return  $T$ 
```

---

Theorems 2 and 3 are the main ingredients of our polynomial algorithm for checking graph weakness (function  $weakOrNotWeak$  in Alg. 1). Our algorithm generates a complete chain of mvs's, that contains at most  $|V|$  mvs's. By Theorem 2, we know that the graph is weak if and only if at some point an mvs with a critical node appears in the chain. Thus, for every new node  $b$  appeared in all such mvs's, it computes a  $b$ -minimal mvs (by backward generating another complete chain of at most  $|V|$  mvs's) and it checks if it is  $b$ -critical. Thanks to Theorem 3, this suffices to conclude.

Function  $immediateMvsRight$  adapts the work in [19] to build a complete chain of mvs's, generated from the bottom mvs  $T = \{s\}$  by iteratively proceeding as follows. Given an mvs  $T$ , pick up any vertex  $u \in T$  and replace it with its immediate successors, denoted by  $succ(u) = \{v \in V : u \rightarrow v \in E\}$ . The set  $S = (T \setminus \{u\}) \cup succ(u)$  is a separator, though not necessarily minimal: some vertices could be covered by other ones. Thus, we consider  $T'' = S \setminus I_t(S)$ , where  $I_t(S)$  contains all the vertices  $v$  such that  $v \sqsubset S \setminus \{v\}$ ; in particular, we always have  $u \in I_t(S)$ . As shown in [19],  $T''$  is an mvs and  $T \sqsubset T''$ .

As an example, let us consider the mvs  $T = \{a, p\}$  in Fig. 3:  $(T \setminus \{a\}) \cup succ(a) = \{b, u, p\}$  is not minimal.  $T^a$  is  $\{u, p\}$ , since the node  $b$  is covered by  $u$ .

The mvs  $T''$  is not necessarily an immediate successor (w.r.t.  $\sqsubset$ ) of  $T$ . To obtain an immediate successor of  $T$  (and thus build a complete chain), it suffices to consider the mvs  $T' = \min_{u \in T} T''$ , where the minimum is calculated w.r.t.  $\sqsubset$ . Indeed, [19] shows that all the immediate successors of an mvs  $T$  can be obtained as  $T''$ , for some  $u \in T$ .

As an example, let us consider again the graph in Fig. 3 and let  $T$  be the mvs  $\{a, b, c\}$ . We have that  $T^a = T^b = \{u, c\} \sqsubset T^c = \{t\}$ .

A  $b$ -minimal predecessor (w.r.t.  $\sqsubset$ ) of  $T$  is computed by function *minimalMvs* in Alg. 2. Given an mvs  $T$  and a node  $u \in T$ ,  $T_u = (T \cup \text{pred}(u)) \setminus I_s(T \cup \text{pred}(u))$  is the analogue of  $T^u$  in the backward direction, where  $\text{pred}(u) = \{v \in V : v \rightarrow u \in E\}$  is the set of immediate predecessors of  $u$ . The correctness of function *minimalMvs* is given by the following result.

**Theorem 4.**  *$T$  is  $b$ -minimal if and only if  $(T \cup \text{pred}(u)) \setminus \{b\} \leq b$ , for every  $u \in T$ .*

**Analysis** First of all, we can calculate the reachability relation for every pair of vertices in  $O(|V|^3)$ , to fill in a  $|V| \times |V|$  binary matrix that allows us to check whether  $T$  is critical in  $O(|V|^2)$  in Alg. 1. Second, relations  $T' \sqsubseteq T$  and  $T \leq T'$  (in *immediateMvsRight* and *minimalMvs*) can be calculated in  $O(|V|^2)$  by calculating the connected component  $C$  of the sink (for  $\sqsubseteq$ ) or of the source (for  $\leq$ ) in the subgraph  $G[V \setminus T]$ ; the desired relation holds if and only if  $T' \cap C = \emptyset$ . Third, given  $T$  and  $u \in T$ , the mvs  $T^u$  can be calculated in  $O(|V|^2)$ , by following [19]: first, calculate the connected component  $C$  of the sink in the subgraph  $G[V \setminus (T \cup \text{succ}(u))]$ ; then,  $I_t(T \cup \text{succ}(u))$  is the subset of  $T \cup \text{succ}(u)$  without an immediate successor in  $C$ . Finally,  $T_u$  in Alg. 2 can be computed symmetrically, by considering the connected component of the source in  $G[V \setminus (T \cup \text{pred}(u))]$  and by excluding all vertices without an immediate predecessor in  $C$ .

Thus, function *immediateMvsRight* costs  $O(|V|^3)$  since it computes at most  $|V|$  mvs's to find  $\min_{u \in T} T^u$ . Function *minimalMvs* in Alg. 2 costs  $O(|V|^4)$ : we need  $O(|V|)$  iterations of the **while** of line 2 and each iteration costs  $O(|V|^3)$ , since, for every  $u \in T$ , we have to check relation  $\leq$ .

Finally, Alg. 1 costs  $O(|V|^5)$ , since function *minimalMvs* is invoked at most once for each node in  $V$ . Indeed, if  $b$  appears as a new node in  $T' \setminus T$  (line 6 of Alg. 1), it cannot have already appeared in a  $T'' \sqsubset T'$  of the chain, otherwise there would exist a walk from  $b$  to  $b$  and hence  $T''$  (as well as  $T'$ ) would be critical and consequently function *weakOrNotWeak* would have terminated in line 5 returning **WEAK** as a result.

## 4 Inefficiency in Flow Networks: Edge-weakness

Interestingly, the notion of weakness strongly depends on the fact that we assign charges to nodes, and not capacities to edges. Indeed, in several settings the two models are interchangeable [1]. However, the analogous of the notion of weakness in the standard flow network model with capacities on edges, that we call *edge-weakness*, does not correspond to weakness.

We first briefly recall some standard notions of flow networks [1]. First, we denote with  $\text{out}(u)$  the set of all edges whose first component is  $u$  and  $\text{in}(u)$  the set of all edges whose second component is  $u$ . A *flow network* is a graph  $G = (V, E)$  endowed with a capacity function  $\{c_e\}_{e \in E}$ , assigning a non-negative number ( $c_e \in \mathbb{R}^+$ ) to every edge. A *flow* in such a network is a function  $f : E \rightarrow \mathbb{R}^+$  such that

- $\forall e \in E. 0 \leq f(e) \leq c_e$ , and
- $\forall u \in V \setminus \{s, t\}. \sum_{e \in \text{in}(u)} f(e) = \sum_{e \in \text{out}(u)} f(e)$ .

The *value* of a flow  $f$ , written  $|f|$ , is defined as  $\sum_{e \in \text{out}(s)} f(e)$  and it turns out to be equal to  $\sum_{e \in \text{in}(t)} f(e)$ .

A flow  $f$  *saturates* a network if, for every path, there exists an edge  $e$  belonging to that path such that  $f(e) = c_e$ . The standard problem in flow networks is to find a saturating flow with maximum value. By mimicking Def. 1, we give the following definition of edge-weak graph:

**Definition 2.** *A graph is edge-weak if there exists a capacity assignment to edges such that the resulting flow network admits a non-maximum saturating flow.*

In general, we can calculate a flow in a network by non-deterministically choosing paths and saturating them. It is easy to see that such an algorithm always calculates a maximum flow if and only if the graph is not edge-weak. In this sense, edge-weakness can be considered another form of inefficiency in network design: to calculate a maximum flow, we cannot use a simple iterated DFS but we need more complex algorithms (see [1]). Something similar happens in [3]. In that model, edges are also equipped with a cost function and the problem is to find a flow of a given value (between 0 and the maximum) but with the lowest possible cost. The authors prove that a greedy algorithm always calculates the minimum cost flow if and only if the underlying graph is series-parallel. Incidentally, we will prove that the notion of series-parallel and non-edge-weakness coincide for acyclic graphs (see Theorem 9 later on).

We now graph-theoretically characterize the notion of edge-weakness. To this aim, recall that a cut of a graph  $G$  is a bipartition of its vertices  $(S, T)$  such that  $s \in S$  and  $t \in T$ ; moreover, its *cut-set* is the set of edges  $(u, v)$  with  $u \in S$  and  $v \in T$ . We call *connected* a cut  $(S, T)$  where every node in  $S$  can be reached from  $s$  without touching nodes in  $T$  and every node in  $T$  can reach  $t$  without touching nodes in  $S$ .

**Theorem 5.**  *$G$  is edge-weak if and only if there exists a walk passing at least twice through the cut-set of some connected cut.*

The proof is similar to the proof of Theorem 1. Just notice that working with connected cuts ensures a minimality property on the associated cut-set: given a connected cut  $(S, T)$ , no cut has a cut-set properly contained in the cut-set of  $(S, T)$ . This somehow corresponds to the minimality condition underlying an mvs.


## 5 Inefficiency in Traffic Networks: Vulnerability

*Traffic Networks* [2,4] provide a model for studying selfish routing: non-cooperative agents travel from a source node  $s$  to a destination node  $t$ . Since the cost (or latency) experienced by an agent while travelling along a path depends on network congestion (and hence on routes chosen by other agents), traffic in a network stabilizes to the equilibrium of a non-cooperative game, where all agents experience the same latency. This phenomenon has been defined by Wardrop [21] in the context of transport analysis.


In the following, we essentially follow the presentation in [17]. Here, a *flow* for a graph  $G = (V, E)$  is a function  $\varphi : P(G) \rightarrow \mathbb{R}^+$ , where  $P(G)$  is the set of paths in  $G$ . A flow induces a unique flow on edges: for any edge  $e \in E$ ,  $\varphi(e) = \sum_{p \in P(G): e \in p} \varphi(p)$ . Since we do not have capacities on edges (as in the standard flow networks) or charges on nodes (as in Depletable Channels), in this model a flow is simply a function assigning non-negative reals to paths, without any further constraint.

A *latency function*  $l_e : \mathbb{R}^+ \rightarrow \mathbb{R}^+$  assigns to each edge  $e$  a latency that depends on the flow on it; as usual, we only consider continuous and non-decreasing latency functions.





**Fig. 5.** The Wheatstone network



**Fig. 6.** Wheatstone optimal subgraph

The latency of a path  $p$  under a flow  $\varphi$  is the sum of the latencies of all edges in the path under  $\varphi$ , i.e.,  $l_p(\varphi) = \sum_{e \in p} l_e(\varphi(e))$ . If  $H$  is a subgraph of  $G$ , we denote with  $l|_H$  the restriction of the latency function  $l$  on the edges of  $H$ .

Given a graph  $G$ , a real number  $r \in \mathbb{R}^+$  and a latency function  $l$ , we call the triple  $(G, r, l)$  an *instance*. A flow  $\varphi$  is *feasible* for  $(G, r, l)$  if the value of  $\varphi$  is  $r$ . Notice that, since we do not have any constraint on edges or vertices, every  $r$  admits at least one feasible flow.

A flow  $\varphi$  feasible for  $(G, r, l)$  is at *Wardrop equilibrium* (or is a *Wardrop flow*) if, for all pairs of paths  $p, q \in P(G)$  such that  $\varphi(p) > 0$ , we have  $l_p(\varphi) \leq l_q(\varphi)$ . In particular, this implies that, if  $\varphi$  is a Wardrop flow, all paths to which  $\varphi$  assigns a positive flow have the same latency. It is known [17] that every instance admits a Wardrop flow and that different Wardrop flows for the same instance have the same latency along the same path. Thus, we denote with  $L(G, r, l)$  the latency of all paths with positive flow at Wardrop equilibrium. In the special case where  $r = 0$ , we let  $L(G, r, l)$  be 0.

*Braess's paradox* [2,5] originates when latency at Wardrop equilibrium decreases because of removing edges (or equivalently, by raising the latency function on edges): an instance  $(G, r, l)$  suffers from Braess's paradox if there is a subgraph of  $G$  with a lower latency. Fig. 5 shows the *Wheatstone network*, a minimal example of Braess's paradox. A Wardrop flow of value 1 assigns all the flow to the path  $suv t$  in the picture. The latency in such a case is 2. In Fig. 6, we show the optimal subgraph: in this case, a Wardrop flow of value 1 assigns  $\frac{1}{2}$  to both paths in the network, thus obtaining a latency of  $\frac{3}{2}$ .

**Definition 3.** A graph  $G$  is *vulnerable* if there exist a value  $r$ , a latency function  $l$  and a subgraph  $H$  of  $G$  such that  $L(G, r, l) > L(H, r, l|_H)$ .

A characterization of vulnerable undirected graphs is presented in [14]. In particular, in [14] it is proved that an undirected graph is vulnerable if and only if it contains (the undirected version of) the Wheatstone network. We now show that the same result holds also for directed graphs, thus answering to Open Question 1 in Sect. 6.1 of [17]. To formally state our result, recall that [20]  $G$  contains a subgraph homeomorphic to  $H$  if  $H$  can be obtained from  $G$  by a sequence of the following operations:

- remove an edge;
- replace  $(u, v)$  and  $(v, w)$  with  $(u, w)$  and delete vertex  $v$ , whenever  $(u, v)$  is the only edge entering into  $v$  and  $(v, w)$  is the only edge leaving  $v$ .

Let us call  $W$  the graph in Fig.1(b), i.e. the graph underlying the Wheatstone network.

**Theorem 6.**  $G$  is vulnerable if and only if it contains a subgraph homeomorphic to  $W$ .

One direction can be proved by assigning a big-enough latency to all edges that do not belong to the homeomorphic copy of  $W$  and by mimicking the latencies in Fig. 5

for the remaining ones. For the converse, we exploit the fact that, for acyclic graphs, vulnerability coincides with not being series-parallel; this is Theorem 1 from [7]<sup>1</sup> and yields the result, by using [8] (where it is proved that, for undirected graphs, being series-parallel coincides with not having a subgraph homeomorphic to  $W$ ; this result scales only to *acyclic* directed graphs). We then reduce the cyclic case to the acyclic one by showing that, in every net that does not contain  $W$ , we can remove cycles without changing the set of acyclic paths and by proving that this operation does not affect vulnerability. Thus, we cannot claim that every non-vulnerable graph is series-parallel but only that it cannot contain a subgraph homeomorphic to  $W$ . An example of a graph that is not vulnerable and not series-parallel will be given in the next section (and, of course, it is cyclic).

To conclude, we should mention [10], where an elegant generalization to all congestion games is given and a characterization of structures that do not suffer of Braess's paradox is given in terms of matroids. Differently from our result, this elegant characterization is not directly related to graph theoretic concepts.

## 6 Comparing Weakness, Edge-Weakness, and Vulnerability

Stemming from the characterizations of weakness (Theorem 1), edge-weakness (Theorem 5) and vulnerability (Theorem 6), we can relate the three notions of inefficiency studied so far. The precise picture (for general directed graphs) is given in the top-left part of Fig. 7.

We first show that vulnerability implies edge-weakness and that it can be characterized by containment of an acyclic weak subgraph.


**Theorem 7.** *If  $G$  is vulnerable, then it is edge-weak.*

**Theorem 8.**  *$G$  is vulnerable if and only if it contains a weak acyclic subgraph.*

It is easy to see that the graph  $W$  is weak, vulnerable and edge-weak. Let us now consider the other graphs in Fig. 7. The graph  $A$  is weak, because its mvs  $\{u\}$  is critical. By contrast, it is not edge-weak (and so neither vulnerable): its only connected cuts are  $(\{s\}, \{u, v, t\})$  and  $(\{s, u, v\}, \{t\})$  and no walk passes through their cut-sets twice. The graph  $B$  is weak (its mvs  $\{u\}$  is critical) and edge-weak (there is a walk passing twice through the cut-set of its connected cut  $(\{s, u\}, \{v, t\})$ ). By contrast, it is not vulnerable, since it does not contain a subgraph homeomorphic to  $W$ . The graph  $C$  is not weak: its mvs's are  $\{s\}$ ,  $\{t\}$ ,  $\{v_1, v_2\}$  and  $\{v_3, v_4\}$  and they are not critical. By contrast, this graph is vulnerable because it trivially contains a subgraph homeomorphic to  $W$ . Finally, let us consider the graph  $D$ : it is not weak (its mvs's are  $\{s\}$ ,  $\{t\}$ ,  $\{u\}$  and  $\{v\}$ , which are not critical) nor vulnerable (it does not contain a subgraph homeomorphic to  $W$ ). By contrast, it is edge-weak, because there is a walk that passes twice through the cut-set of the connected cut  $(\{s, u, x\}, \{y, v, t\})$ .

Graph  $A$  also testifies that the characterization of vulnerability given in [7] does hold only for irredundant graphs. Indeed, Lemma 2 in [7] does not hold for  $A$  (that is

<sup>1</sup> Indeed, acyclic graphs are a specific case of what they call *irredundant* graphs. A graph is said to be irredundant if every edge and every node belongs to a simple (i.e., acyclic)  $st$ -path.



**Fig. 7.** The inclusion diagram for cyclic (top-left) and acyclic (top-right) directed graphs ( $W$  is the Wheatstone graph – see Fig. 1(b))

redundant because of the vertex  $v$ ): the graph is cyclic but it does not contain what they call an  $s$ - $t$  paradox (see Def. 5 in [7]).

If we restrict ourselves to *acyclic* graphs (i.e., DAGs), the inclusion diagram changes; it is depicted in the top-right part of Fig. 7. Indeed, for DAGs, we can prove that both vulnerability and edge-weakness coincide with not being *two-terminal series-parallel* (TTSP [16]); furthermore, the resulting class (properly) contains weak DAGs.

**Theorem 9.** *If  $G$  is a DAG, the following statements are equivalent:*

1.  $G$  is vulnerable;
2.  $G$  is not TTSP;
3.  $G$  is edge-weak;
4.  $G$  contains a weak subgraph.

## 7 Conclusion

We have studied different models for networks: one is provided with depletable node charge modeling energy consumption in ad-hoc networks; one is the standard flow network model; and one is provided with latency functions on the edges in a game theoretic framework for modeling traffic networks. In all models, a graph-theoretical notion of inefficiency can be identified, always related to a counterintuitive phenomenon: some networks increase their performances when an edge is removed. We have compared these forms of inefficiency and we have found precise relationships.

We have also shown a polynomial time algorithm for checking weakness. Our algorithm may seem overly complicated. Indeed, because of Theorem 3, it would be enough to compute, for every  $b \in V$ , a  $b$ -minimal mvs and check whether it is critical or not. The problem is that not every vertex occurs in an mvs (see, for example, vertex  $v$  of graph A in Fig. 7). Moreover, for a vertex  $b$  occurring in an mvs, we have not found an efficient way to directly compute a  $b$ -minimal mvs; indeed, Alg. 2 computes a  $b$ -minimal mvs starting from an mvs that already contains  $b$ . Efficiently finding such an mvs is left for future development.

The characterization we gave for vulnerability gives also hints on its computational complexity: it is polynomial for general graphs (a  $O(|V|^5)$  algorithm can be easily extracted from the proof of Theorem 6). Finally, edge-weakness is polynomial in the

acyclic case (since it coincides with vulnerability) but we still do not know in the general case; possibly, by following the ideas underlying the algorithm for weakness, a polynomial-time algorithm for edge-weakness may be devised.

**Acknowledgements** The authors wish to thank Irene Finocchi and Fabrizio Grandoni for helpful discussions about the topic of this paper.

## References

1. R. Ahuja, T. Magnanti, and J. Orlin. *Network Flows, theory, algorithms, and applications*. Prentice-Hall, New Jersey, 1993.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. W. Bein, P. Brucker, and A. Tamir. Minimum cost flow algorithms for series-parallel networks. *Discrete Applied Mathematics*, 10:117–124, 1985.
4. M. Bell and Y. Iida. *Transportation Network Analysis*. Wiley, 1987.
5. D. Braess. Über ein paradoxon aus der verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.
6. P. Cenciarelli, D. Gorla, and I. Salvo. Depletable channels: Dynamics and behaviour. In *Proc. of FCT09*, volume 5690 of *LNCS*, pages 50–61. Springer, 2009.
7. X. Chen, Z. Diao, and X. Hu. Excluding braess paradox in nonatomic selfish routing. In *Proc. of SAGT15*, volume 9347 of *LNCS*, pages 219–230. Springer, 2015.
8. R. J. Duffin. Topology of series-parallel networks. *Journ. of Math Analysis and Applications*, 10:303–318, 1965.
9. F. Escalante. Schnittverbände in graphen. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 38(1):199–220, 1972.
10. S. Fujishige, M. Goemans, T. Harks, B. Peis, and R. Zenklusen. Matroids are immune to braess paradox. Available at: [arXiv:1504.07545](https://arxiv.org/abs/1504.07545), 2015.
11. J. Golumbic. *Algorithmic Graph Theory and Perfect Graphs*. Academic Press, 1980.
12. M. Iri. Theory of uncontrollable flows – a new type of network-flow theory as a model for the 21<sup>st</sup> century of multiple values. *Computers Math. Applic.*, 35(10):107–123, 1998.
13. T. Kloks and D. Kratsch. Listing all minimal separators of a graph. *SIAM J. Comput.*, 27(3):605–613, 1998.
14. I. Milchtaich. Network topology and the efficiency of equilibrium. *Games and Economic Behavior*, 57:321–346, 2006.
15. G. Newell. *Traffic Flow on Transportation Networks*. MIT Press, 1980.
16. J. Riordan and C. Shannon. The number of two-terminal series-parallel networks. *Journal of Mathematics and Physics*, 21:83–93, 1942.
17. T. Roughgarden. On the severity of Braess’s Paradox: Designing networks for selfish users is hard. *J. Comput. Syst. Sci.*, 72(5):922–953, 2006.
18. Y. Sheffi. *Urban Transportation Networks*. Prentice-Hall, 1985.
19. H. Shen and W. Liang. Efficient enumeration of all minimal separators in a graph. *Theor. Comput. Sci.*, 180(1-2):169–180, 1997.
20. J. Valdes, R. Tarjan, and E. Lawler. The recognition of series-parallel digraphs. *SIAM Journal of Computing*, 11:298–313, 1982.
21. J. Wardrop. Some theoretical aspects of road traffic research. In *Proc. of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.
22. B. Warneke, M. Last, B. Liebowitz, and K. S. J. Pister. Smart dust: Communicating with a cubic-millimeter computer. *IEEE Computer*, 34(1):44–51, 2001.

## A Proofs

### A.1 Proofs of Section 2

#### Proof of Theorem 1

(If) We define  $\eta(v) = 1$  for all  $v \in T$  and we let  $\eta(v)$  be sufficiently large on all other vertices, so that  $\max_\eta = |T|$ . We may assume without loss of generality that the walk  $r : a \rightsquigarrow b$  is acyclic and such that  $r \cap T = \{a, b\}$ . Similarly, by the minimality of  $T$ , there exist (acyclic) directed walks  $p : s \rightsquigarrow a$  and  $q : b \rightsquigarrow t$  such that  $p \cap T = \{a\}$  and  $q \cap T = \{b\}$ . Then, a flow  $\phi$  of one unit along the path  $s \rightsquigarrow a \rightsquigarrow b \rightsquigarrow t$  is feasible and it leaves the channel with a charge-to-node assignment  $\theta$  such that  $\max_\theta \leq |T| - 2$ . Then, we can combine  $\phi$  with a maximum flow for  $\theta$  to obtain a dead network after a flow of value  $\max_\theta + 1 < \max_\eta$ . Thus,  $G$  is weak.

(Only if) Let  $\eta$  be a channel inhibited by some flow  $\phi$  of value  $n < \max_\eta$  and call  $\zeta$  the resulting (dead) channel. There exists an mvs  $T$  in the graph such that  $\zeta(v) = 0$ , for all  $v \in T$ . Since every flow cannot exceed the capacity of an mvs,<sup>2</sup>  $\max_\eta \leq \sum_{v:v \in T} \eta(v)$ . By definition  $\zeta(v) = \eta(v) - \sum_{p:v \in p} \phi(p)$  and hence  $\eta(v) = \sum_{p:v \in p} \phi(p)$ . Suppose no directed walk exists between any two vertices of  $T$ . Thus, since  $T$  is an mvs, all paths must include *precisely* one vertex in  $T$ ; hence,  $\sum_{v:v \in T} \sum_{p:v \in p} \phi(p) = \sum_p \phi(p)$ . Summing up, we have the absurd:

$$\max_\eta \leq \sum_{v:v \in T} \eta(v) = \sum_{v:v \in T} \sum_{p:v \in p} \phi(p) = \sum_p \phi(p) = n. \quad \square$$

### A.2 Proofs of Section 3

**Proof of Theorem 2** Looking for a contradiction, let us suppose that in a graph there is a critical mvs  $T$  with a return in  $b \in T$ , but there exists a complete chain  $T_0, T_1, \dots, T_n$  such that for all  $i$ ,  $T_i$  does not contain any critical node.

Trivially,  $b \not\subseteq T_0$  and  $b \subseteq T_n$  (this happens for every  $b \in V \setminus \{s\}$ ); thus, let us consider the index  $i$  such that  $b \not\subseteq T_i$  and  $b \subseteq T_{i+1}$ . By construction,  $T_{i+1} = T_i^{u_i}$ , for some  $u_i \in T_i$ ; moreover,  $b \notin T_{i+1}$ , because, by hypothesis, no critical node belongs to any mvs in the chain.

<sup>2</sup> This comes from the min-cut-max-flow theorem [1], that can be easily rephrased in our setting to sound as *min-mvs-max-flow*. Indeed, we can adopt the standard translation from vertex-capacitated nets into edge-capacitated nets: we replace every vertex  $v$  different from  $s$  and  $t$  with two new vertices  $v_i$  and  $v_o$  and we add a new edge  $v_i \rightarrow v_o$  whose capacity is the charge of node  $v$ ; every edge  $u \rightarrow v$  is replaced with the edge  $u_o \rightarrow v_i$  whose capacity is  $\infty$ ; every edge  $s \rightarrow v$  is replaced with the edge  $s \rightarrow v_i$  whose capacity is  $\infty$ ; every edge  $u \rightarrow t$  is replaced with the edge  $u_o \rightarrow t$  whose capacity is  $\infty$ . Now, there is a one-to-one correspondence between the mvs's of the vertex-capacitated model and the cuts of the corresponding edge-capacitated model whose cut-set is formed only by edges  $v_i \rightarrow v_o$ ; moreover, the capacity of any mvs and of the corresponding cut coincide. Since only cuts arising from mvs's have a finite capacity, it is easy to show that a min-cut comes from a min-mvs and, conversely, that a min-mvs induces a min-cut. Moreover, since only edges of the form  $v_i \rightarrow v_o$  put constraints on a flow, a max-flow in the edge-capacitated model corresponds to a max-flow in the vertex-capacitated model, and vice versa.

Let us define the many-steps predecessors of  $b$  as  $\text{pred}^*(b) = \{v \in V : v \rightsquigarrow b\}$  and take  $P = \text{pred}^*(b) \cap T_i$ , the set of many-steps predecessors of  $b$  in  $T_i$ . For this set, we observe two things:

- (i)  $\emptyset \neq P \leq b$ : since  $T_i$  is an mvs, every path passing through  $b$  must cross  $T_i$  but, since  $b \not\sqsubseteq T_i$ , it must be that  $T_i \leq b$ . Trivially, since all nodes between  $s$  and  $b$  are by definition the many-step predecessors of  $b$ , the same relation holds by restricting  $T_i$  to such nodes, thus obtaining  $P$ . Since  $b \neq s$ , this entails that  $P \neq \emptyset$ .
- (ii) For all  $p \in P$ ,  $T_i^p = T_{i+1}$ : If some  $p \in P$  belonged to  $T_{i+1}$ , we could find a walk  $p \rightsquigarrow b \rightsquigarrow u$  with  $u \in T_{i+1}$  (indeed,  $b \sqsubseteq T_{i+1}$ ); this would make  $T_{i+1}$  a critical mvs. So, we must have that all nodes in  $P$  disappear in  $T_{i+1}$ ; by construction, this happens because, for every  $p \in P$ , we have that  $p \in I_t(T_i \cup \text{succ}(u_i))$  and hence  $p \sqsubseteq T_{i+1}$ . This fact, together with  $p \notin T_{i+1}$ , entails that  $T_i^p \sqsubseteq T_{i+1}$ , for all  $p \in P$ . If  $T_i^p \neq T_{i+1}$  for some  $p \in P$ , we contradict the hypothesis that  $T_0, T_1, \dots, T_n$  is a complete chain: indeed, we would have that  $T_i \sqsubset T_p \sqsubset T_{i+1}$ .

Because of point (ii), the set  $U$  of new nodes added from  $T_i$  to  $T_{i+1}$  is the same whenever we refine  $T_i$  with  $u_i$  or any other  $p \in P$ ; in particular, every node in  $U$  is an immediate successor of every node in  $P$ . Moreover, we can claim the following fact about  $U$ :

- (iii)  $b \sqsubseteq U$ : since  $b \sqsubseteq T_{i+1}$ , every walk from  $b$  to  $t$  must cross  $T_{i+1}$ . If it passes through a node  $x \in T_i \setminus T_{i+1}$ , then the walk  $p \rightsquigarrow b \rightsquigarrow x$ , for any  $p \in P$ , would make  $T_i$  critical.

We now use these facts to contradict the assumption that  $b$  is a critical node. Consider all the paths of the form  $s \rightsquigarrow p \rightarrow u \rightsquigarrow t$ , with  $p \in P$  and  $u \in U$ . Every mvs, to cut such paths, must contain either a set of nodes  $P' \leq P$  or a set of nodes  $U' \sqsupseteq U$ . In both cases, since  $P' \leq b$  and  $b \sqsubseteq U'$  (because of (i) and (iii)),  $b$  cannot belong to any mvs (see Lemma 1) and hence cannot be a critical node.  $\square$

**Proof of Theorem 3** Let  $T$  be a  $b$ -critical mvs, where  $a \in T$  is such that  $a \rightsquigarrow b$ . By contradiction, assume the existence of an mvs  $T^*$  that is both  $b$ -minimal and not  $b$ -critical. We shall now prove that this will entail that  $a$  and  $b$  cannot both belong together to the same mvs, thus contradicting the existence of  $T$ .

Clearly,  $a \notin T^*$ , because  $T^*$  is not  $b$ -critical. Moreover, the existence of a walk from the source  $s$  to  $a$  that passes through  $T^*$  would imply the existence of a walk starting in  $T^*$  that reaches  $a$  and then  $b$ , by contradicting that  $T^*$  is not  $b$ -critical. Therefore, it must be  $a \sqsubseteq T^*$ .

Let  $A = \{u \in T^* \mid a \rightsquigarrow u \neq b\} = \{a_1, \dots, a_n\}$ . Notice that  $A \neq \emptyset$ , otherwise we would have  $a \sqsubseteq b$  and therefore  $a$  and  $b$  could not belong to the same mvs (see Lemma 1). For each element  $a_i \in A$ , let us now consider the mvs  $T_{a_i}^* \sqsubseteq T^*$ . By  $b$ -minimality of  $T^*$ ,  $b$  does not belong to  $T_{a_i}^*$ . This implies that there exists a set of nodes  $P_i \subseteq \text{pred}(a_i)$  such that  $P_i \leq b$ . Observe that there is some path from  $s$  to  $b$  that do not pass through  $a$ , otherwise  $a \leq b$  and then they could not belong to the same mvs. Therefore, for each  $i \in \{1, \dots, n\}$ , we can consider the set of nodes  $B_i \subseteq P_i$  that are on a path from  $s$  to  $b$  that does not touch  $a$ . Notice that, for all  $i$ , we have that  $B_i \neq \emptyset$ , otherwise all paths from  $s$  to  $b$  would pass through  $a$ .

Let us now fix an  $i \in \{1, \dots, n\}$ , let us come back to  $T$  and consider how it can cut all paths of the form  $s \rightsquigarrow b_i \rightarrow a_i \rightsquigarrow t$ , for all  $b_i \in B_i$ .

If in  $T$  there exists a set of nodes  $L_i \leq B_i$ , this would imply that  $b$  cannot be in  $T$ , because  $L_i \cup \{a\} \leq B_i \cup \{a\} \leq b$  and  $L_i \cup \{a\} \subseteq T$  (see Lemma 1).

So, it must be that at least one path  $s \rightsquigarrow b_i \rightarrow a_i$  does not pass through  $T$ . Therefore there must exist a set in  $R_i \subseteq T$  such that  $a_i \sqsubseteq R_i$ . Since this argument works for every  $i$ , we can consider  $R = \bigcup_i R_i \subseteq T$ . Then,  $A = \{a_1, \dots, a_n\} \sqsubseteq \bigcup_i R_i = R$ . But again, since  $a \sqsubseteq A \cup \{b\} \sqsubseteq R \cup \{b\}$  and  $R \cup \{b\} \subseteq T$ , Lemma 1 would imply that  $a$  cannot belong to  $T$ . This is a contradiction with the initial choice of the mvs  $T$  that contains both  $a$  and  $b$ .  $\square$

**Proof of Theorem 4** Given an mvs  $T$  and two nodes  $u, b \in T$ , let us consider the mvs  $T_u \sqsubset T$ . We have that  $b \notin T_u$  if and only if  $b \in I_s(T \cup \text{pred}(u))$ , i.e.  $(T \cup \text{pred}(u)) \setminus \{b\} < b$ . Since every predecessor (w.r.t.  $\sqsubset$ ) of  $T$  can be obtained as  $T_u$ , for some  $u \in T$ , we easily conclude.  $\square$

### A.3 Proofs of Section 4

Let us denote with  $\text{cutset}(S, T)$  the cut-set of the cut  $(S, T)$ .

**Lemma 2.** *Let  $(S, T)$  be a connected cut. For every cut  $(S', T')$ , it holds that  $\text{cutset}(S', T') \not\subseteq \text{cutset}(S, T)$ .*

*Proof.* By contradiction, assume a cut  $(S', T')$  such that  $\text{cutset}(S', T') \subset \text{cutset}(S, T)$  and let  $(u, v) \in \text{cutset}(S, T) \setminus \text{cutset}(S', T')$ . Because  $(S', T')$  is a cut and  $(u, v) \notin \text{cutset}(S', T')$ , it can either be  $\{u, v\} \subseteq S'$  or  $\{u, v\} \subseteq T'$ .

In the first case, since  $(S, T)$  is a connected cut, we know that there exists a walk  $v \rightsquigarrow t$  containing only vertices of  $T$ . But  $v \in S'$  whereas  $t \in T'$ ; hence, there must exist a  $(x, y) \in v \rightsquigarrow t$  such that  $x \in S'$  and  $y \in T'$ . Then,  $(x, y) \in \text{cutset}(S', T')$ , whereas  $(x, y) \notin \text{cutset}(S, T)$ , because  $\{x, y\} \subseteq T$ . This contradicts the assumption  $\text{cutset}(S', T') \subset \text{cutset}(S, T)$ .

In the second case, we work in a similar way, but consider the walk  $s \rightsquigarrow u$  containing only vertices of  $S$ .  $\square$

**Proof of Theorem 5** Given a set of edges  $X$ , we write  $c(X)$  to denote  $\sum_{x \in X} c(x)$ .

(If) Let  $(S, T)$  be the connected cut and  $E'$  be its cut-set. Let  $p$  be a walk that passes through  $E'$  at least twice, with  $(u, v)$  and  $(x, y)$  be the first and the last edge of  $E'$  touched by it. We define  $c_e$  as the number of occurrences of  $e$  in  $p$ , for all  $e \in E'$ , and we let  $c_e$  be  $c(E') + 1$  on all other edges.

First of all, notice that the maximum flow has value  $c(E')$ . This follows from the min-cut-max-flow theorem [1], by noting that  $(S, T)$  is a minimum cut. Indeed, by Lemma 2, no cut has a cut-set contained in  $E'$ ; moreover, by definition of  $c$ , any cut whose cut-set contains an edge not belonging to  $E'$  has a capacity greater than  $c(E')$ .

By definition of cut-set,  $u \in S$  and  $y \in T$ ; moreover, by construction,  $s \rightsquigarrow u$  and  $y \rightsquigarrow t$  do not pass through  $E'$ . Then, a flow  $f$  of one unit along  $p$  is feasible and it leaves the net with a residual capacity lower than  $c(E') - 1$ . Then, we can combine  $f$  with a

maximum flow for the residual net to obtain a saturating flow of value at most  $c(E') - 1$ . Thus,  $G$  is edge-weak.

(Only If) Let  $c$  be a capacity function that admits a saturating flow  $f$  of value smaller than the maximum and call  $c'$  the resulting residual capacity. Since  $f$  saturates  $c$ , we now show that there exists a connected cut whose cut-set has a residual capacity that equals 0. If this was not the case, let us reason as follow. Start with the cut  $(\{s\}, V \setminus \{s\})$ . Clearly, this is a connected cut; so, there exists an edge  $e_1$  from  $s$  to some  $u_1 \in V \setminus \{s\}$  such that  $c'_{e_1} > 0$ . If  $u_1 = t$ , we have a contradiction with the fact that  $f$  saturates  $c$ . So, consider the cut  $(\{s, u_1\}, V \setminus \{s, u_1\})$ . Again, this is a connected cut and so there exists an edge  $e_2$  from  $\{s, u_1\}$  to some  $u_2 \in V \setminus \{s, u_1\}$  such that  $c'_{e_2} > 0$ . If  $u_2 = t$ , we have a contradiction with the fact that  $f$  saturates  $c$ . Otherwise, we go on, until we find a  $u_k = t$ . Then, the path  $s \rightarrow u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_k = t$  contradicts the fact that  $f$  saturates  $c$ .

To conclude, let  $(S, T)$  be the connected cut whose cut-set has been saturated by  $f$ . By the min-cut-max-flow theorem [1],  $c(E') \geq |f_{\max}|$ . But this is possible only if there is a path that passes through  $E'$  at least twice. Indeed, if this was not the case, we would have that  $c(E') = |f|$ , in contradiction with the assumption  $|f| < |f_{\max}|$ .  $\square$

#### A.4 Proofs of Section 5

First, we characterize vulnerability for acyclic graphs; this is an immediate corollary of the main result from [7]: an irredundant graph is vulnerable if and only if it is not series-parallel. Since acyclic graphs are a special case of irredundant ones, the theorem needs no proof.

**Theorem 10.** *Let  $G$  be acyclic; then,  $G$  is vulnerable if and only if it is not TTSP.*

We are now ready to move to general (i.e., cyclic) graphs. To this aim, let  $G$  be a directed  $st$ -graph; we denote with  $\text{Path}_A(G)$  the set of acyclic paths from  $s$  to  $t$  in  $G$ .


**Lemma 3.** *Let  $G$  be a directed  $st$ -graph that does not contain a subgraph homeomorphic to the Wheatstone network. Let  $C = v_1, \dots, v_n$  be the nodes of a simple cycle in  $G$ . Then, we can remove an edge in  $C$ , obtaining a graph  $G'$  such that  $\text{Path}_A(G') = \text{Path}_A(G)$ .*

*Proof.* Let  $(v_1, v_2), (v_2, v_3), \dots, (v_{n-1}, v_n), (v_n, v_1)$  be a simple cycle in  $G$ . Without loss of generality, let  $v_1$  be a node such that there exists a walk  $s \rightsquigarrow v_1$  that touches  $C$  only in  $v_1$ . Such a node does exist, otherwise the cycle would not be reachable from the source  $s$ . Similarly, let  $v_k$  be the last node in the sequence  $v_1, \dots, v_n$  such that there exists a walk  $v_k \rightsquigarrow t$  that touches  $C$  only in  $v_k$ . Such a node does exist, otherwise the sink  $t$  would not be reachable from the cycle  $C$ .

If  $k = 1$ , no acyclic path in  $G$  passes through an edge in  $C$  (see Fig. 8(a)). In such a case, we can remove all edges in  $C$ , without removing any acyclic path in  $G$ .

Let us consider now a walk  $v_1 \rightsquigarrow t$ . If all such walks have the form  $v_1 \rightsquigarrow v_k \rightsquigarrow t$ , then  $v_k$  is the only exit from  $C$ . Consequently, the edge  $(v_k, v_{k+1 \bmod n})$  does not belong to any acyclic path in  $G$ . Indeed all paths of the form  $s \rightsquigarrow v_k \rightarrow v_{k+1 \bmod n} \rightsquigarrow t$  must return in  $v_k$  after  $v_{k+1 \bmod n}$  to leave the cycle  $C$ .





**Fig. 8.** Figures for Lemma 3

Otherwise, let  $v_i$  ( $1 \leq i < k$ ) be the first node in  $v_1, \dots, v_{k-1}$  that can reach  $t$  without passing through other nodes in  $C$  and consider nodes  $v_{i+1}, \dots, v_k$ . If all walks from  $s$  to them pass through other nodes in  $C$ , again the edge  $(v_k, v_{k+1 \bmod n})$  does not belong to any acyclic path in  $G$ . Indeed, consider a path that uses edge  $(v_k, v_{k+1 \bmod n})$ . For what we have just assumed, such a path enters into the cycle  $C$  in a node  $v' \in \{v_{k+1 \bmod n}, \dots, v_1, \dots, v_i\}$  and has to pass through edge  $(v_k, v_{k+1 \bmod n})$ . Since all vertices in  $v_{k+1 \bmod n}, \dots, v_1, \dots, v_i$  can reach  $t$  only through  $v_i$  or after it, the path is cyclic (see Fig. 8(b)).

Otherwise, let  $v_j$  ( $i < j \leq k$ ) be such that there exists a walk  $s \rightsquigarrow v_j$  that does not pass through other nodes in  $C$ . In this case,  $G$  would contain a subgraph homeomorphic to the Wheatstone network, as given in Fig. 8(c).  $\square$

**Lemma 4.** Let  $G' \subseteq G$ . Let  $\varphi$  be an (acyclic) flow at the Wardrop equilibrium for  $(G, r, l)$ . If  $\text{Path}_A(G') = \text{Path}_A(G)$  then  $\varphi$  is a flow at the Wardrop equilibrium for  $(G', r, l)$ .

*Proof.* By Proposition 2.2 in [17], we know that, for every  $p, q \in P(G)$ , if  $\varphi_p > 0$  then  $\ell_p(\varphi) \leq \ell_q(\varphi)$ . By hypothesis,  $\varphi$  assigns positive flow only to acyclic paths in  $G$ ; thus,  $p \in \text{Path}_A(G) = \text{Path}_A(G')$ . Moreover, since  $G' \subseteq G$ , it holds that  $P(G') \subseteq P(G)$ . Thus, trivially, for every  $p, q \in P(G')$ , if  $\varphi_p > 0$  then  $\ell_p(\varphi) \leq \ell_q(\varphi)$ . Again by Proposition 2.2 in [17], this means that  $\varphi$  is a flow at the Wardrop equilibrium for  $(G', r, l)$ .

**Lemma 5.** Let  $G' \subseteq G$ . If  $\text{Path}_A(G') = \text{Path}_A(G)$  then  $G$  is vulnerable if and only if  $G'$  is vulnerable.

*Proof.* Vulnerability of  $G'$  trivially entails vulnerability of  $G$ . Let us prove the opposite implication. Let  $H \subset G$  be such that  $L(H, r, l) < L(G, r, l)$ , for some  $r$  and  $l$ .

If  $H \subset G'$ , let  $\varphi$  be an acyclic flow for  $G$  at the Wardrop equilibrium (one always exists by Proposition 2.4 of [17]). By Lemma 4,  $\varphi$  is a flow for  $G'$  at the Wardrop equilibrium; thus,  $L(H, r, l) < L(G', r, l)$ , i.e.  $G'$  is vulnerable.


Otherwise, it cannot be  $H = G'$ , because  $\text{Path}_A(H) \subset \text{Path}_A(G)$ ; indeed, because of Lemma 4, if  $\text{Path}_A(H) = \text{Path}_A(G)$ , we would have  $L(H, r, l) = L(G, r, l)$ . Thus, there is an edge  $(u, v) \in H$  such that  $(u, v) \notin G'$ ; this means that  $(u, v)$  only belongs to cyclic paths of  $G$ , because by hypothesis  $\text{Path}_A(G') = \text{Path}_A(G)$ . We now show that this implies the existence of a simple cycle  $p : w \rightsquigarrow w$  in  $G$  containing  $(u, v)$  such that:

1. for every  $w' \in w \rightsquigarrow u$ , it holds that  $w \sqsubseteq w'$ , and
2. for every  $w' \in v \rightsquigarrow w$ , it holds that  $w' \leq w$ .

First of all, since  $(u, v)$  belongs to a cyclic path in  $G$ , it also belongs to a simple cycle in  $G$ , say  $w \rightsquigarrow w$ . To prove the first claim, consider  $w' \neq w$  (the case for  $w' = w$  is trivial). Then, if there was a path  $s \rightsquigarrow w'$  not passing through  $w$ , then  $s \rightsquigarrow w' \rightsquigarrow u \rightarrow v \rightsquigarrow w \rightsquigarrow t$  would be an acyclic path in  $G$ , again by the assumption that  $(u, v)$  only belongs to cyclic paths in  $G$ . Similarly, to prove the second claim, assume a walk  $w' \rightsquigarrow t$  not passing through  $w$ ; then  $s \rightsquigarrow w \rightsquigarrow u \rightarrow v \rightsquigarrow w' \rightsquigarrow t$  would be an acyclic path in  $G$ , again the assumption that  $(u, v)$  only belongs to cyclic paths in  $G$ .

Then, consider  $H' = H \setminus p$ . Trivially,  $\text{Path}_A(H') = \text{Path}_A(H)$  and let  $\varphi'$  be an acyclic flow for  $H$  at the Wardrop equilibrium (one always exists by Proposition 2.4 of [17]). By Lemma 4,  $\varphi'$  is a flow for  $H'$  at the Wardrop equilibrium and  $L(H, r, l) = L(H', r, l)$ . If  $H' \subset G'$ , then  $L(H', r, l) < L(G', r, l)$  and  $G'$  is vulnerable. Otherwise, we can find another simple cycle to be removed from  $H'$  but this procedure has to terminate, eventually yielding that  $G'$  is vulnerable, as desired.  $\square$

**Proof of Theorem 6** For the "if" part, we know that  $G$  admits a subgraph of the form



Let us consider the latency assignment  $l$  that assigns:


- 0 to all edges in  $u \rightsquigarrow v$ ;
- $x$  to the first edge in  $s' \rightsquigarrow u$  and in  $v \rightsquigarrow t'$  and 0 to all the remaining edges in those paths;
- 1 to the first edge in  $s' \rightsquigarrow v$  and in  $u \rightsquigarrow t'$  and 0 to all the remaining edges in those paths;
- $\infty$  to all the remaining edges.

Trivially, this reproduces the Wheatstone network within  $G$  that, consequently, is vulnerable.

For the "only if" part, if  $G$  is not cyclic, the statement follows by Theorem 10. If  $G$  contains cycles, by contradiction, let us suppose that it does not contain a subgraph homeomorphic to the Wheatstone network. Then, by Lemma 3, we can transform  $G$  into an acyclic subgraph  $G'$  such that  $\text{Path}_A(G') = \text{Path}_A(G)$ . By Lemma 5, since  $G$  is vulnerable,  $G'$  is vulnerable too; thus, since  $G'$  is acyclic and not TTSP (by Theorem 10), it contains a subgraph homeomorphic to the Wheatstone network [8] and, consequently, also  $G$  does. Absurd.  $\square$

## A.5 Proofs of Section 6

**Proof of Theorem 7** By Theorem 6,  $G$  admits a subgraph of the form




Let us consider the capacity assignment  $\{c_e\}_{e \in E}$  that assigns

- 2 to all edges in  $s \rightsquigarrow s'$  and  $t' \rightsquigarrow t$ ;
- 1 to all edges in  $s' \rightsquigarrow u \cup s' \rightsquigarrow v \cup u \rightsquigarrow v \cup u \rightsquigarrow t' \cup v \rightsquigarrow t'$ ; and
- 0 to all the remaining ones.


Trivially, the flow assigning 1 to the path  $s \rightsquigarrow s' \rightsquigarrow u \rightsquigarrow v \rightsquigarrow t' \rightsquigarrow t$  is a non-maximum saturating flow, since there exists a flow with value 2.  $\square$

**Proof of Theorem 8** If  $G$  is vulnerable, then, by Theorem 6, it admits a subgraph  $H$  of the form:



The subgraph  $H$  is weak. As a matter of fact,  $\{u, v\}$  is an mvs for  $H$ , and the walk  $u \rightsquigarrow v$  allows us to conclude by Theorem 1.

For the converse implication, let  $H$  be a weak subgraph of  $G$ . Theorem 1 implies that  $H$  admits an mvs  $T$  such that there exists a walk  $a \rightsquigarrow b$  in  $H$ , with  $\{a, b\} \subseteq T$ . By minimality of  $T$ , there exist in  $H$  a walk from  $s$  to  $a$  (that does not contain  $b$ ), a walk from  $s$  to  $b$  (that does not contain  $a$ ), a walk from  $a$  to  $t$  (that does not contain  $b$ ), and a walk from  $b$  to  $t$  (that does not contain  $a$ ). Thus, we have found the following subgraph of  $H$  (and, hence, also of  $G$ ):



Notice that both  $s'$  and  $t'$  cannot occur in  $a \rightsquigarrow b$ , otherwise we would have a cycle in  $H$ , in contradiction with the assumption of its acyclicity. For the same reason,  $a \neq b$ . Then, we can easily conclude that  $G$  is vulnerable by Theorem 6.  $\square$

### Proof of Theorem 9

(1)  $\Rightarrow$  (2): By Theorem 7.

(2)  $\Rightarrow$  (3): We prove the contrapositive and work by induction on the structure of  $G$ .

The base case is when  $G$  is a single edge  $(s, t)$ . In this case, the only saturating flow for  $G$  is the one that saturates the capacity of  $(s, t)$ ; trivially, this is the maximum flow.

If  $G$  is the serial composition of  $G_1$  and  $G_2$  (that are TTSP), by the inductive hypothesis  $G_i$  is not edge-weak, for  $i \in \{1, 2\}$ . Now, since  $E = E_1 \cup E_2$ , every capacity assignment  $\{c_e\}_{e \in E}$  is  $\{c_e\}_{e \in E_1} \cup \{c_e\}_{e \in E_2}$ . Moreover, because of serial composition,  $\max(G, \{c_e\}_{e \in E}) = \min\{\max(G_1, \{c_e\}_{e \in E_1}), \max(G_2, \{c_e\}_{e \in E_2})\}$ . Finally, since every source-to-sink path in  $G$  is a source-to-sink path in  $G_1$  followed by a source-to-sink path in  $G_2$ , every flow  $f$  for  $G$  induces a flow with the same value  $f_1$  for  $G_1$  and

a flow with the same value  $f_2$  for  $G_2$ . Now, by contradiction, assume that there exists  $\{c_e\}_{e \in E}$  such that  $(G, \{c_e\}_{e \in E})$  is edge-weak, i.e. it admits a saturating flow  $\hat{f}$  such that  $|\hat{f}| < \max(G, \{c_e\}_{e \in E})$ . Thus,  $\hat{f}_i$  saturates all the paths of  $G_i$ , for either  $i = 1$  or  $i = 2$ . But this is in contradiction with the inductive hypothesis, since  $|\hat{f}_i| < \max(G_i, \{c_e\}_{e \in E_i})$ .

If  $G$  is the parallel composition of  $G_1$  and  $G_2$  (that are TTSP), the proof is similar. Just notice that in this case  $\max(G, \{c_e\}_{e \in E}) = \max(G_1, \{c_e\}_{e \in E_1}) + \max(G_2, \{c_e\}_{e \in E_2})$ , that every source-to-sink path in  $G$  is a source-to-sink path either in  $G_1$  or in  $G_2$  and, consequently, that every flow  $f$  for  $G$  induces two flows,  $f_1$  for  $G_1$  and  $f_2$  for  $G_2$ , such that  $|f| = |f_1| + |f_2|$ . By contradiction, let  $\{c_e\}_{e \in E}$  be such that  $(G, \{c_e\}_{e \in E})$  admits a saturating flow  $\hat{f}$  with  $|\hat{f}| < \max(G, \{c_e\}_{e \in E})$ . Then,  $\hat{f}_i$  saturates all the paths of  $G_i$ , for  $i \in \{1, 2\}$ . The contradiction comes from the fact that  $|\hat{f}_1| + |\hat{f}_2| < \max(G_1, \{c_e\}_{e \in E_1}) + \max(G_2, \{c_e\}_{e \in E_2})$  and from  $|\hat{f}_i| \leq \max(G_i, \{c_e\}_{e \in E_i})$ , for both  $i \in \{1, 2\}$ .

- (3)  $\Rightarrow$  (4): If  $G$  is a not TTSP DAG, then, by [8], it admits a subgraph homeomorphic to the Wheatstone graph; such a subgraph is weak (see the proof of Theorem 8).  
(4)  $\Rightarrow$  (1): By Theorem 8. □